



**Sandia
National
Laboratories**

Characterizing Resource Utilization on Production HPC Platforms to Inform Sustainable Policy, Procurements, and Development

Douglas M. Pase, Anthony M. Agelastos, Gary Lawson, Ron A. Oldfield, Joel O. Stevenson, Scott A. Warnock

February 2020

Sandia National Laboratories is a multission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

ABSTRACT

The high performance computing (HPC) ecosystem consists of four chief community members: application developers (developers), analysts who apply the applications for programmatic needs (users), system administrators/support (sysadmins), and platform designers/procurers (designers). This ecosystem thrives when the community member's needs are aligned and understood. There is a disconnect at times between analysts' actual platform needs and each of the different community member's perception of the analysts' needs. This disconnect *may* cause sub-optimal HPC platform policies (i.e., from the sysadmins), future platforms that don't maximize throughput (i.e., from the designers), and algorithms not used as intended/tuned (i.e., from developers). The research discussed with this project aims to characterize HPC platform utilization by users and to disseminate that information to developers, sysadmins, and designers to maintain alignment and health of the SNL HPC ecosystem. Some questions being answered and quantified that would benefit the community are enumerated below. A longer-term goal of this project is to make this information readily available and part of SNL's HPC community's processes.

1. What is the distribution of jobs and platform resource usage (i.e., node hours) for jobs of a given size (i.e., node count)?
2. How do jobs use the available memory?
3. How many jobs fail (i.e., what percentage) and how much of the cluster resources (i.e., node hours) does that represent?
4. How well do requested job time limits reflect actual time usage?
5. Are users making reasonable use of the available resources (e.g., number of cores and network) or are they using only a small number of cores out of those available?
6. Are users making use of available Hyperthreads, or is that a wasted resource?